arXiv:1511.06733v2 [stat.ME] 16 Jun 2016

On an inferential model construction using generalized associations

Ryan Martin
Department of Mathematics, Statistics, and Computer Science
University of Illinois at Chicago
rgmartin@uic.edu

June 17, 2016

Abstract

The inferential model (IM) approach, like fiducial and its generalizations, depends on a representation of the data-generating process. Here, a particular variation on the IM construction is considered, one based on generalized associations. The resulting generalized IM is more flexible than the basic IM in that it does not require a complete specification of the data-generating process and is provably valid under mild conditions. Computation and marginalization strategies are discussed, and two applications of this generalized IM approach are presented.

Keywords and phrases: Likelihood; marginalization; Monte Carlo; plausibility function; random set; validity.

1 Introduction

An advantageous feature of the mainstream approaches to statistical inference is simplicity. On one hand, likelihood-based approaches, including "Frasian" inference (e.g., Barndorff-Nielsen 1991; Fraser 1990, 1991, 2011; Reid 2003) and certain forms of Bayesian inference (e.g., Berger et al. 2009, 2015; Bernardo 1979; Ghosh 2011), are simple in the sense that the calculations relevant to data analysis are largely (or completely) determined by the posited sampling model. On the other hand, frequentist approaches are also simple because the "do whatever works well" viewpoint is extremely flexible. This is in sharp contrast with fiducial inference (Barnard 1995; Dawid and Stone 1982; Fisher 1973; Taraldsen and Lindqvist 2013), its generalizations (Hannig 2009; Hannig et al. 2015), and the recently proposed inferential model (IM) framework (Martin and Liu 2013, 2015a,b,c), which appear to be not-so-simple in the sense that their construction depends on something more than the data and sampling model. In particular, the fiducial and IM construction begins with a specific representation of the data-generating mechanism, one that determines but is not determined by the sampling model. This data-generating mechanism identifies an auxiliary variable, or pivotal quantity, that controls the random variation in the observable data. A familiar example of this kind is the regression model,

$Y = X\beta + \sigma\varepsilon$, where the random "$\varepsilon$" part controls the variation of the response Y around the deterministic "$X\beta$" part. That the fiducial and IM solutions depend on the choice of the data-generating mechanism may be seen as a shortcoming of these approaches.

One approach to deal with the choice of data-generating mechanism is to find one that is "best" in some sense; for example, Pal Majumdar and Hannig (2015) compare different data-generating mechanisms using higher-order asymptotics in the fiducial context. Since defining and identifying the "best" is difficult, I want to take a different approach. In this paper, building on Martin and Liu (2015b, Ch. 11), I want to incorporate the familiar frequentists' flexibility into the IM construction. This allows the user to construct a *generalized IM* without specifying a full data-generating mechanism, simplifying the construction in several ways. First, just like in the likelihood-based approaches mentioned above, a generalized IM can be constructed based on the sampling model alone, or some function thereof, easing the burden on the user. Second, the generalized IM can be constructed based on a *generalized association* that involves only a one-dimensional auxiliary variable, which simplifies user's task of selecting a good predictive random set. Compare this to the basic IM approach where the user must first specify a data-generating mechanism and carry out some potentially non-trivial dimension-reduction steps (e.g., Martin and Liu 2015a). Despite making substantial simplifications to the IM construction, it can be shown that this generalization preserves the IM's guaranteed validity property under mild conditions. Therefore, the generalized IM framework is a simple and widely applicable tool for valid, prior-free, probabilistic inference.

This paper's main contribution is the new perspective it brings to some more-or-less familiar ideas, results, and techniques. Specifically, all of the familiar considerations used in constructing statistical procedures fit within the the seemingly rigid IM framework, and this has at least two useful consequences. First, working within the IM framework does not require that one abandon all the classical tools and ways of thinking—these can be merged seamlessly into the framework itself. Second, new insights concerning these classical tools can be gained when looking from an IM point of view; see Section 3.3.

The remainder of the paper is organized as follows. After some background on IMs in Section 2, the new generalized IM approach is presented in Section 3, with a motivating validity theorem and a special case that is relatively easy to implement, involving only a scalar auxiliary variable, and having good properties. Some important practical considerations, namely, computation and marginalization, are discussed in Section 4, and two interesting and challenging applications—inference on the odds ratio in 2×2 tables and inference on the error variance in mixed-effects models—are presented in Section 5. Concluding remarks are made in Section 6.

2 Background on IMs

Let $Y \in \mathbb{Y}$ be the observable data, and write $\mathsf{P}_{Y|\theta}$ for the sampling model, which depends on an unknown parameter $\theta \in \Theta$. In the basic IM framework, described in Martin and Liu (2013), the starting point—the *A-step*—is to associate Y and θ with an unobservable auxiliary variable $U \in \mathbb{U}$ with known distribution P_U. Formally, suppose the association can be written as

$$Y = a(\theta, U), \quad U \sim \mathsf{P}_U. \tag{1}$$

Martin and Liu (2015a,c) argue that some dimension-reduction steps should be taken first before an association mapping is defined, so the left-hand side may be something different than the observable data, e.g., a minimal sufficient statistic. This dimension-reduction step is recommended, but it is not necessary to describe these details here. The result of the A-step is a set-valued mapping

$$\Theta_y(u) = \{\theta : y = a(\theta, u)\}, \quad u \in \mathbb{U}, \tag{2}$$

indexed by the observed $Y = y$. The main point is that the association determines the sampling model $\mathsf{P}_{Y|\theta}$ or, alternatively, the ingredients in (1) must be chosen to be consistent with the given sampling model. However, there may be several versions of the association that are consistent with the sampling model, and different versions may produce different inferences. This is not unlike the frequentists' choice of (approximate) pivot for constructing a test, confidence region, etc. In any case, the question of which association (1) to take, for given sampling model $\mathsf{P}_{Y|\theta}$, is an important one.

The second step in the basic IM construction—the *P-step*—is to predict the unobserved value of U in (1), corresponding to the observed $Y = y$, with predictive random set $\mathcal{S}$. The P-step is the defining feature of the IM framework, driving its essential properties and separating it from the approach described in Dempster (2008). The distribution $\mathsf{P}_{\mathcal{S}}$ of $\mathcal{S}$ is to be chosen by the user, subject to a certain "validity" condition, namely, that, if $f_{\mathcal{S}}(u) = \mathsf{P}_{\mathcal{S}}(\mathcal{S} \ni u)$, then

$$f_{\mathcal{S}}(U) \geq_{\text{st}} \mathsf{Unif}(0,1), \quad \text{as a function of } U \sim \mathsf{P}_U,$$

where "$\geq_{\text{st}} \mathsf{Unif}(0,1)$" means "stochastically no smaller than $\mathsf{Unif}(0,1)$," i.e.,

$$\mathsf{P}_U\big\{f_{\mathcal{S}}(U) \leq \alpha\big\} \leq \alpha, \quad \forall\, \alpha \in (0,1). \tag{3}$$

Intuitively, the random set $\mathcal{S}$ is meant to be "good" at predicting samples from P_U and (3) makes this precise: the $\mathsf{P}_{\mathcal{S}}$-probability of the event "$\mathcal{S} \ni u$" is small only for a set of u values with relatively small P_U-probability. Sufficient conditions for (3) are mild, so it is easy to find a valid predictive random set; in fact, most applications of IMs employ a simple "default" predictive random set, see (13).

The third and final step in the basic IM construction—the *C-step*—is to combine the association at the observed data $Y = y$ with the predictive random set $\mathcal{S}$. Specifically, one obtains a random subset of Θ:

$$\Theta_y(\mathcal{S}) = \bigcup_{u \in \mathcal{S}} \Theta_y(u). \tag{4}$$

The intuition behind this is as follows: if one believes that $\mathcal{S}$ contains the value of U corresponding to the observed $Y = y$ and the true θ, which is justified by (3), then one must also believe, with equal conviction, that $\Theta_y(\mathcal{S})$ contains the true θ. The IM output is the distribution of the random set $\Theta_y(\mathcal{S})$, which I will summarize with a plausibility function. Specifically, if $A \subset \Theta$, then the plausibility function at A is

$$\mathsf{pl}_y(A) = \mathsf{P}_{\mathcal{S}}\{\Theta_y(\mathcal{S}) \cap A \neq \varnothing\}.$$

Of course, the plausibility function depend on $\mathcal{S}$ or, more precisely, on $\mathsf{P}_{\mathcal{S}}$, but I omit this dependence in the notation. For interpretation, $\mathsf{pl}_y(A)$ is a measure of the degree of

belief, given data y, in the falsity of "$\theta \notin A$." The user's "belief" is first encoded in $\mathsf{P}_{\mathcal{S}}$, a personal or belief probability, subject to the constraint (3), which is then transferred to the parameter space in the IM's C-step. Intuitively, it is possible that two disjoint assertions are highly plausible based on the given data, and the plausibility function allows for this, i.e., plausibility satisfies $\mathsf{pl}_y(A) + \mathsf{pl}_y(A^c) \geq 1$ for all A. Moreover, Theorem 2 in Martin and Liu (2013) shows that if $\mathcal{S}$ satisfies (3), then $\mathsf{pl}_Y(A)$ is properly calibrated as a function of $Y \sim \mathsf{P}_{Y|\theta}$ for fixed A, in the sense that

$$\sup_{\theta \in A} \mathsf{P}_{Y|\theta}\big\{\mathsf{pl}_Y(A) \leq \alpha\big\} \leq \alpha, \quad \forall\, \alpha \in (0,1), \quad \forall\, A \subseteq \Theta, \tag{5}$$

or, in other words, for any $A \subseteq \Theta$, if $\theta \in A$, then $\mathsf{pl}_Y(A) \geq_{\text{st}} \mathsf{Unif}(0,1)$, as a function of $Y \sim \mathsf{P}_{Y|\theta}$. When (5) holds, the IM is said to be *valid*. This validity property aids in interpreting the plausibility function values—it puts the personal/belief probabilities on an objective $\mathsf{Unif}(0,1)$ scale—and also facilitates the construction of IM-based decision rules with guaranteed error rate control.

The conclusion I hope the reader will reach from this brief summary is that the IM approach is conceptually straightforward and accomplishes what Fisher's fiducial approach was meant to, namely, valid prior-free probabilistic inference. The apparent cost is that the IM output depends on the choice of association (1), the choice of predictive random set, and, in a less-obvious way, on the dimension of the auxiliary variable. The need to specify an association, carry out the necessary dimension-reduction steps, and introduce a valid predictive random set may give the impression that the IM approach is not user-friendly. The goal of this paper is to show how one can construct a valid IM by dealing with these challenges indirectly.

3 A class of generalized IMs

3.1 Construction

Towards accomplishing the goals laid out above, we discuss here how the basic association (1) can be made simpler and more flexible, by relaxing the direct connection with the sampling model and informally reducing auxiliary variable dimension, while still retaining the desirable validity properties of the resulting IM.

Start by going back to the beginning of Section 2 where only the sampling model $\mathsf{P}_{Y|\theta}$ for data Y given parameter θ is available. The IM construction in Section 2 is based on identification of an unobservable auxiliary variable U to associate with (Y, θ) and then to be predicted. The basic approach identifies U by thinking about the data-generating process, but this is potentially restrictive and unnecessary. Rather than specifying a potentially relatively high-dimensional auxiliary variable corresponding to a data-generation process, and then subsequently reducing the dimension according to guidelines in Martin and Liu (2015a,c), is it possible to specify an auxiliary variable of the appropriate dimension directly and easily?

Towards answering this question, the key insight is that the association in (1) need not involve the full data Y. For a function $(y, \theta) \mapsto T_{y,\theta}$, consider a *generalized association*

$$T_{Y,\theta} = a(\theta, U), \quad U \sim \mathsf{P}_U, \tag{6}$$

where U is some auxiliary variable taking values in a space $\mathbb{U}$. Note that, unless $y \mapsto T_{y,\theta}$ is one-to-one for each θ, which is not a useful case, the generalized association does not determine the sampling model for Y, thereby relaxing the requirement in Section 2 that the association specify a version of the data-generating mechanism. It does, however, determine the sampling model of $T_{Y,\theta}$ under $Y \sim \mathsf{P}_{Y|\theta}$, so (6) is compatible with $\mathsf{P}_{Y|\theta}$ in this sense. The function $T_{Y,\theta}$ can depend on θ or not, and its distribution need not be continuous. Some examples are discussed below and in the later sections.

Based on (6), the (generalized) A-step defines the set-valued mapping

$$\Theta_y(u) = \{\theta : T_{y,\theta} = a(\theta, u)\}, \quad (y, u) \in \mathbb{Y} \times \mathbb{U}. \tag{7}$$

Then the P- and C-steps can be carried out exactly like in Section 2. In particular, the P-step introduces a valid random set $\mathcal{S} \sim \mathsf{P}_{\mathcal{S}}$ for predicting the unobserved value of U in (6), and the C-step yields the random set $\Theta_y(\mathcal{S})$ as in (4) and the corresponding belief and plausibility functions bel_y and pl_y, depending implicitly on $\mathsf{P}_{\mathcal{S}}$. I will call the resulting IM a *generalized IM* and, interestingly, validity of this generalized IM, in the sense of (5), follows immediately from the construction.

Theorem 1. *For the generalized association* (6)*, let $\mathcal{S} \sim \mathsf{P}_{\mathcal{S}}$ be a valid predictive random set for $U \sim \mathsf{P}_U$. If $\Theta_y(\mathcal{S}) \neq \varnothing$ with $\mathsf{P}_{\mathcal{S}}$-probability 1 for all y, then the generalized IM is valid in the sense of* (5)*.*

Proof. For any A, let (y, θ, u) be such that $\theta \in A$ and $T_{y,\theta} = a(\theta, u)$. Since $\{\theta\} \subset A$ and pl_y is monotone, we have

$$\mathsf{pl}_y(A) \geq \mathsf{pl}_y(\{\theta\}) = \mathsf{P}_{\mathcal{S}}\{\Theta_y(\mathcal{S}) \ni \theta\} = \mathsf{P}_{\mathcal{S}}\{\mathcal{S} \ni u\} = f_{\mathcal{S}}(u).$$

Since $f_{\mathcal{S}}(U) \geq_{\text{st}} \mathsf{Unif}(0,1)$ as a function of $U \sim \mathsf{P}_U$, it follows that $\mathsf{pl}_Y(A) \geq_{\text{st}} \mathsf{Unif}(0,1)$ as a function of $Y \sim \mathsf{P}_{Y|\theta}$, i.e., $\mathsf{P}_{Y|\theta}\{\mathsf{pl}_Y(A) \leq \alpha\} \leq \alpha$, for all $\alpha \in (0,1)$. This holds for all $\theta \in A$, so take supremum of the left-hand side over $\theta \in A$ to complete the proof. □

Therefore, construction of a valid generalized IM is possible and seems to be fairly straightforward. An important consequence of the validity theorem is that plausibility regions based on the generalized IM have the nominal coverage probability. That is, if

$$\mathcal{C}_\alpha(y) = \{\theta : \mathsf{pl}_y(\theta) > \alpha\},$$

where $\mathsf{pl}_y(\theta) = \mathsf{pl}_y(\{\theta\})$, then

$$\mathsf{P}_{Y|\theta}\{\mathcal{C}_\alpha(Y) \ni \theta\} = \mathsf{P}_{Y|\theta}\{\mathsf{pl}_Y(\theta) > \alpha\},$$

and since the validity property (5) holds for all A, in particular, $A = \{\theta\}$, we get that the right-hand side in the above display is $\geq 1 - \alpha$ for all θ. An important observation is that this does not require large samples or any assumptions on the model.

3.2 A useful special case

There are, of course, a variety of ways one can specify the generalized association (6). Here I will elaborate on one simple but general strategy. Let $(y, \theta) \mapsto T_{y,\theta}$ be scalar-valued, e.g., the likelihood ratio or a function thereof; in general, it is not a statistic

because it depends on θ. Moreover, since the map is scalar-valued, in most cases, it cannot be one-to-one so it corresponds to a non-trivial summary of the data y. Suppose that $T_{Y,\theta}$ has a continuous distribution, under $Y \sim \mathsf{P}_{Y|\theta}$, and let F_θ be the corresponding distribution function. Now specify an association in terms of the distribution of $T_{Y,\theta}$:

$$T_{Y,\theta} = F_\theta^{-1}(U), \quad U \sim \mathsf{Unif}(0,1). \tag{8}$$

The case of discrete $T_{Y,\theta}$ can be handled similarly, i.e.,

$$F_\theta(T_{Y,\theta}-) \leq U < F_\theta(T_{Y,\theta}), \quad U \sim \mathsf{Unif}(0,1),$$

where $F_\theta(t-) = \lim_{s \uparrow t} F_\theta(s)$ is the left-hand limit. This corresponds to taking $a(\theta, u)$ in (6) to be $F_\theta^{-1}(u)$. Now, with a suitable predictive random set for $U \sim \mathsf{Unif}(0,1)$, this generalized association leads to a valid generalized IM.

Corollary 1. *The generalized IM constructed based on the association* (8) *and a valid predictive random set $\mathcal{S}$ for $U \sim \mathsf{Unif}(0,1)$ is valid in the sense of Theorem 1, provided that $\Theta_y(\mathcal{S}) \neq \varnothing$ with $\mathsf{P}_{\mathcal{S}}$-probability 1 for all y.*

This provides a simple and general procedure for constructing a valid generalized IM based on a choice of mapping $T_{y,\theta}$. In fact, this shows that the work done by Martin (2015) in the frequentist context is just a special case of the proposed generalized IM framework. His choice to work primarily with the negative log-likelihood ratio,

$$T_{y,\theta} = -2 \log \frac{L_y(\theta)}{\sup_{\vartheta \in \Theta} L_y(\vartheta)}, \tag{9}$$

with L_y the likelihood function for θ based on data $Y = y$. The $T_{y,\theta}$ in(9) is the deviance used frequently in Schweder and Hjort (2016). Also, other authors, e.g., Wasserman (1990), Aickin (2000), and Denœux (2014), have used the likelihood ratio to construct a plausibility function for statistical inference, but in a different way than I propose here. There are, however, other choices of $T_{y,\theta}$; see, e.g., Remark 3 and Section 5.

A natural question is if anything is gained from the generalized IM perspective, besides the apparent simplicity, compared to the basic IM approach described in Section 2 and the references therein. The next example demonstrates that the simple generalized IM can lead to improved efficiency, at least in some cases.

Example 1. Let $Y_1, \ldots, Y_n$ be iid samples from a $\mathsf{Gamma}(\theta_1, \theta_2)$ distribution, where θ_1 is the shape parameter and θ_2 is the scale parameter, both unknown. This same problem was considered Martin and Liu (2015a, Section 5.3) and they presented a basic IM solution based on a reduction to the complete sufficient statistic. This requires specifying a predictive random set for a two-dimensional auxiliary variable consisting of two independent uniforms. No IM optimality results are available for this problem, so they made the natural choice of a square-shaped predictive random set. This guarantees validity of the IM, but efficiency is a question. For comparison, consider a generalized IM based on the likelihood ratio, which is also valid; the computational details are discussed in Section 4.1. I simulate $n = 25$ observations from the gamma distribution with $\theta_1 = 7$ and $\theta_2 = 3$. Figure 1 displays several results: the Jeffreys prior Bayesian posterior samples, the 90% confidence ellipse based on asymptotic normality of the maximum likelihood

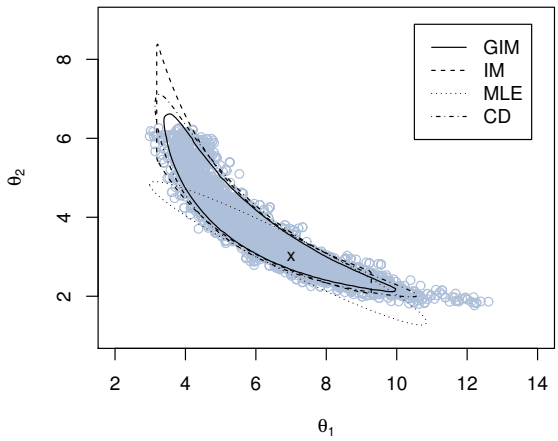

Figure 1: Output from the gamma simulation in Example 1: Jeffreys prior Bayes posterior samples (gray); maximum likelihood-based 90% confidence ellipse (dotted); 90% deviance-based confidence region (dot-dashed); and 90% plausibility regions based on basic (dashed) and generalized (solid) IMs.

estimator, the 90% confidence region based on the asymptotic chi-square distribution of the deviance (Schweder and Hjort 2016), the 90% plausibility region based on the IM construction in Martin and Liu (2015a), and the 90% plausibility region based on the likelihood ratio-based generalized IM. Interestingly, the generalized IM plausibility region has guaranteed 90% coverage and it captures the overall shape of the posterior, which is non-elliptical. It is also slightly smaller than the deviance-based region and is considerably smaller than the basic IM plausibility region. Another exact confidence region for this gamma problem is obtained in Taraldsen and Lindqvist (2013).

3.3 Remarks

Remark 1 (on asymptotics). To make this discussion concrete, consider the case where Y consists of a collection of n iid observations. When n is large, there is no shortage of pivotal quantities $T_{Y,\theta}$ that can be used in the generalized association (8). Indeed, Wilks's theorem says that $T_{Y,\theta}$ in (9) has an asymptotic chi-square distribution under $\mathsf{P}_{Y|\theta}$, as $n \to \infty$. In this case the F_θ in (8) can, asymptotically, be taken as a suitable chi-square distribution function, free of θ. The same holds in the case with nuisance parameters using a profile likelihood, as in (12). There are many other choices of $T_{Y,\theta}$ that are asymptotic pivots, e.g., the quantities in Brazzale et al. (2007, Chap. 8) with higher-order approximation accuracy. The point is that the generalized IM framework provides a tool for valid statistical inference without appealing to asymptotics but, if desired, asymptotic theory can be used just to provide simple large-sample approximations.

Remark 2 (on confidence distributions). Confidence distributions (Schweder and Hjort 2002, 2016; Singh et al. 2007; Xie and Singh 2013) have received considerable attention recently, especially in the meta-analysis context (Claggett et al. 2014; Liu et al. 2015, 2014; Xie et al. 2011; Yang et al. 2014), a primary selling point being that it "unifies" (Xie and Singh 2013, p. 3) existing approaches. Their point is that a variety of standard tools can

be converted into a confidence distribution or an asymptotic confidence distribution. My proposal here for a generalized IM can be interpreted similarly, since many familiar ideas from classical statistics can be employed to construct a valid generalized IM.

Remark 3 (on efficiency and choice of $T_{Y,\theta}$). Towards an optimal IM, Martin and Liu (2013) suggested that, for a fixed θ_0, the best random set $\mathcal{S}$ is one that makes $\mathsf{pl}_Y(\theta_0)$ as stochastically as small as possible, subject to the validity condition. They argue that there exists a nested collection $\mathbb{Y}_\alpha \subset \mathbb{Y}$, depending on $\mathcal{S}$ and θ_0, such that $\mathsf{pl}_y(\theta_0) > \alpha$ if and only if $y \in \mathbb{Y}_\alpha$ and, furthermore, the optimal $\mathcal{S}$ has corresponding $\mathbb{Y}_\alpha$ such that

$$\int_{\mathbb{Y}_\alpha} S_\theta(y)\, p_\theta(y)\, dy = 0 \quad \text{at } \theta = \theta_0 \text{ for all } \alpha,$$

where p_θ is the density function for Y and $S_\theta(y) = (\partial/\partial\theta) \log p_\theta(y)$ is the familiar score function. Since $\mathsf{E}_{\theta_0}\{S_{\theta_0}(Y)\} = 0$, this condition implies that $\mathbb{Y}_\alpha$ is suitably balanced with respect to the distribution of $S_{\theta_0}(Y)$; this is called a *score-balance condition*. A set that will satisfy the score-balance condition, at least asymptotically, is

$$\mathbb{Y}_\alpha = \{y : S_{\theta_0}(y)^\top I(\theta_0)^{-1} S_{\theta_0}(y) \leq c_\alpha\}$$

for suitable constant c_α, where $I(\theta)$ is the Fisher information. This suggests choosing

$$T_{Y,\theta_0} = S_{\theta_0}(Y)^\top I(\theta_0)^{-1} S_{\theta_0}(Y),$$

and the corresponding plausibility function matches (asymptotically) the p-value of Rao's score test, which has certain optimality properties. This provides some insight into the choice of an efficient mapping $T_{y,\theta}$, but more work is needed. How the optimality considerations in the confidence distribution context (e.g., Schweder and Hjort 2016, Ch. 5) might be useful in the IM context deserves further investigation.

4 Practical considerations

4.1 Computation

For the case (8), suppose that large values of $T_{y,\theta}$ are suggestive that the model $\mathsf{P}_{Y|\theta}$ does not fit data $Y = y$ well. The log-likelihood ratio in (9), the score-balanced cased in Remark 3, among others, are of this form. In this case, a natural choice of the random set $\mathcal{S}$ is the one-sided (nested) random interval

$$\mathcal{S} = [0, U], \quad U \sim \mathsf{Unif}(0, 1).$$

With this choice,

$$\begin{aligned} \Theta_y(\mathcal{S}) \cap A \neq \varnothing &\iff \{\theta : F_\theta(T_{y,\theta}) \leq U\} \cap A \neq \varnothing \\ &\iff \{U \geq F_\theta(T_{y,\theta}),\ \exists\, \theta \in A\} \\ &\iff \Big\{U \geq \inf_{\theta \in A} F_\theta(T_{y,\theta})\Big\} \end{aligned}$$

and, therefore, the corresponding plausibility function is

$$\mathsf{pl}_y(A) = \mathsf{P}_{\mathcal{S}}\{\Theta_y(\mathcal{S}) \cap A \neq \varnothing\} = 1 - \inf_{\theta \in A} F_\theta(T_{y,\theta}) = \sup_{\theta \in A} \bar{F}_\theta(T_{y,\theta}), \tag{10}$$

where $\bar{F}_\theta = 1 - F_\theta$ is the survival function. Of course, for singleton assertions, no optimization is necessary. The point is that evaluating the generalized IM plausibility function requires only some relatively simple probability calculations.

In cases where the distribution function F_θ is not available in closed form, a conceptually simple Monte Carlo approximation is available:

$$F_\theta(t) \approx \frac{1}{M} \sum_{m=1}^{M} I\{T_{Y^{(m)},\theta} \leq t\}, \tag{11}$$

where $\{Y^{(m)} : m = 1, \ldots, M\}$ are independent copies of $Y \sim \mathsf{P}_{Y|\theta}$. Of course, if direct information about the distribution of $T_{Y,\theta}$ is available, e.g., that it depends only on some function of Y, then this can be used to avoid simulation of the entire Y. This approach is straightforward, but can be time-consuming to implement because the plausibility function may need to be evaluated at many different θ values, and each requires its own Monte Carlo simulation. This difficulty can be avoided if it were possible to simulate from $\mathsf{P}_{Y|\theta}$ for only a single value of θ. One way this can be achieved is if it happens that the distribution of $T_{Y,\theta}$, under $Y \sim \mathsf{P}_{Y|\theta}$, does not depend on θ, i.e., $F_\theta \equiv F$. This invariance property holds if $T_{Y,\theta}$ is itself a pivot, which can be arranged in some examples (e.g., Martin 2015, Sec. 2.4). More generally, an importance sampling strategy can be employed to approximate F_θ over a range of θ with only a single Monte Carlo sample. Choose a fixed parameter value, say, $\hat{\theta}$, a suitable estimator, and rewrite (11) as

$$F_\theta(t) \approx \frac{1}{M} \sum_{m=1}^{M} I\{T_{Y^{(m)},\theta} \leq t\} \frac{L_{Y^{(m)}}(\theta)}{L_{Y^{(m)}}(\hat{\theta})},$$

where, this time, $\{Y^{(m)} : m = 1, \ldots, M\}$ are independent samples from $\mathsf{P}_{Y|\hat{\theta}}$, and can be reused for different values of θ. This is reminiscent of parametric bootstrap (e.g., Davison and Hinkley 1997), and will have a much smaller computational cost compared to the naive Monte Carlo approximation in (11). The two Monte Carlo strategies discussed here are extreme in the sense that the former takes a Monte Carlo sample for each θ while the latter takes only one Monte Carlo sample for a single $\hat{\theta}$. Various middle ground strategies are also possible, e.g., take Monte Carlo samples for a fixed grid $\{\vartheta_1, \ldots, \vartheta_G\}$ of parameter values and do an importance sampling-based approximation of $F_\theta(t)$ using samples corresponding to grid point ϑ_g where $g = \arg\min_h \|\theta - \vartheta_h\|$.

Example 2. An interesting non-standard example is the so-called asymmetric triangular distribution (e.g., Berger et al. 2009, Example 11), with density function

$$p_\theta(y) = \begin{cases} 2y/\theta & \text{if } 0 \leq y \leq \theta, \\ 2(1-y)/(1-\theta) & \text{if } \theta < y \leq 1, \end{cases}$$

where $\theta \in [0, 1]$. The density has a unique mode at θ, but the density has a corner and is not differentiable there. Consider making inference on θ based on an independent sample

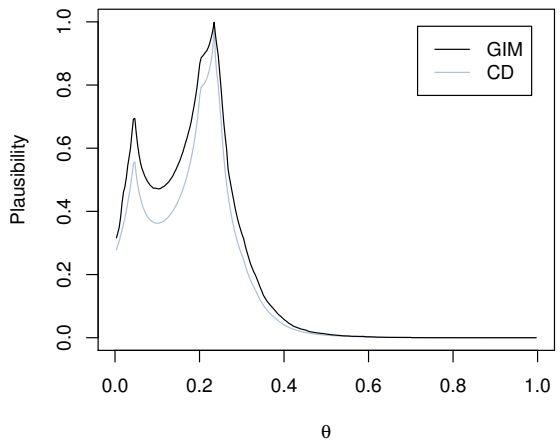

Figure 2: Plots of the plausibility function $\mathsf{pl}_y(\theta)$ for the triangular model in Example 2 based on naive (dashed) and importance sampling-driven (solid) Monte Carlo.

$Y = (Y_1, \ldots, Y_n)$. This is a challenging problem because there is no non-trivial sufficient statistic and the formal Fisher information is not well-defined. Constructing an efficient IM using the basic approach outlined in Section 2 is difficult because there is no clear strategy to reduce the dimension of the auxiliary variable. However, a generalized IM for θ is readily available here using the likelihood ratio (9) as in Section 3.2. For a quick comparison of the generalized IM (based on importance sampling) and the confidence distribution based on the asymptotic chi-square distribution of the deviance in Schweder and Hjort (2016), data Y of size $n = 10$ is simulated from the triangular distribution with $\theta = 0.3$. Plots of the plausibility and confidence curves are shown in Figure 2. The two curves have roughly the same shape, though the confidence curve is a bit tighter, a consequence of the overly optimistic asymptotic approximation.

Though the context here is a bit different, the use of Monte Carlo methods to construct tests and confidence regions has been addressed previously in the literature. For example, our $100(1-\alpha)\%$ plausibility regions correspond to finding solutions to the equation $\mathsf{pl}_y(\theta) = \alpha$. When the plausibility function can only be evaluated via Monte Carlo, solving this equation is a stochastic approximation problem (Robbins and Monro 1951), and has been discussed in Garthwaite and Buckland (1992) and Botev and Lloyd (2015); see, also, Bølviken and Skovlund (1996).

Another issue to address is optimization of the function $F_\theta(T_{y,\theta})$ over a subset A of θ values. This will again be relevant in the discussion of marginalization below. Recently, but again in a slightly different context, Xiong (2015) considers this optimization problem and suggests some localization strategies as well as a proper choice of grid points, via space-filling designs, on which the plausibility function surface can be built up.

4.2 Handling nuisance parameters

Most practical problems involve nuisance parameters, so having some general techniques to eliminate these parameters is important. Without loss of generality, partition the full parameter θ as $\theta = (\psi, \lambda)$, where ψ is the interest parameter and λ is the nuisance

parameter. Here I will discuss three different approaches for eliminating λ in to construct a marginal generalized IM for ψ.

A first strategy is conditioning. In particular, let $y \mapsto (T_y, T_y')$ be a one-to-one transformation of y, independent of the parameter θ. If the conditional distribution of T_y, given T_y', is free of λ, then this conditional distribution can be used to construct a generalized IM for ψ. Section 5.1 presents an example of this conditioning strategy in action.

The second strategy is a direct marginalization by selecting a function $T_{Y,\psi}$, depending on Y and ψ only, such that its distribution is free of the nuisance parameter λ. A general candidate for such a function, generalizing the idea at the end of Section 3, is the profile likelihood ratio

$$T_{Y,\psi} = -2 \log \frac{\sup_\lambda L_y(\psi, \lambda)}{\sup_{\psi,\lambda} L_y(\psi, \lambda)}. \tag{12}$$

Composite transformation models (Barndorff-Nielsen 1988) form a general class of problems where this approach to marginalization can be applied. For example, in the two-parameter gamma model, where ψ is the shape, the Bartlett test statistic has distribution free of the nuisance scale parameter. Similarly, in the bivariate normal model, where ψ is the correlation, the sample correlation coefficient $\hat{\psi}$ has distribution free of λ, the means and variances; the profile likelihood is a function of only $\hat{\psi}$ and ψ and, therefore, also has distribution free of λ. A mixed-effects model, where the nuisance fixed-effect parameters are eliminated via marginalization, is presented in Section 5.2.

A third strategy, which seems to be unique to the framework presented here, is a different form of marginalization via optimization. When the underlying random sets are nested, which is the recommended choice, the plausibility function is called *consonant* (e.g., Shafer 1987). In particular, this means that the plausibility function evaluated at a set A equals the suprema of the plausibility function evaluated at points in A. This provides some further explanation for the expression for $\mathsf{pl}_y(A)$ in (10) involving a supremum. This is relevant in the present situation because a problem that involves nuisance parameters can be handled by considering assertions about the full parameter (ψ, λ) that span the full range of λ. Therefore, marginalization can be accomplished by optimization *after* evaluating the plausibility function, compared to the pre-plausibility evaluation optimization in the profiling approach discussed above. This further demonstrates the importance of the optimization aspects discussed in Section 4.1. It is preferable to eliminate the nuisance parameters before evaluating plausibility, if possible, because it reduces the computational cost, but for some problems there are no obvious conditioning or profiling strategies to use, so this default marginalization tool is necessary.

5 Applications

5.1 Odds ratio in a 2×2 table

Let $Y = (Y_0, Y_1)$ be two independent binomial counts, with $Y_0 \sim \mathsf{Bin}(n_0, \theta_0)$ and $Y_1 \sim \mathsf{Bin}(n_1, \theta_1)$, where $n = (n_0, n_1)$ is known but $\theta = (\theta_0, \theta_1)$ is unknown. Data such as these arise in, say, a clinical trial, where Y_0 and Y_1 correspond to the number of events observed

under the control and treatment. Suppose that the quantity of interest is the odds ratio

$$\psi = \frac{\theta_1/(1-\theta_1)}{\theta_0/(1-\theta_0)}.$$

As in Hannig and Xie (2012), a key observation is that the conditional distribution of Y_1, given $Y_0 + Y_1$, depends on ψ only, not on the nuisance parameter θ_0 (or θ_1), though the distribution form is not a standard one. In particular,

$$\mathsf{P}(Y_1 = y_1 \mid Y_0 + Y_1 = t) \propto \binom{n_1}{y_1}\binom{n_0}{t-y_1}\psi^{y_1},$$

with y_1 ranging over $\max\{t - n_0, 0\}$ and $\min\{n_1, t\}$. As discussed in Section 4.2, let $T_Y = Y_1$ and $T_Y' = Y_0 + Y_1$. For the observed value t of T_Y', let $F_{t,\psi}$ be the conditional distribution function corresponding to the mass function in the above display. The resulting generalized association is

$$F_{t,\psi}(Y_1 - 1) \leq U < F_{t,\psi}(Y_1), \quad U \sim \mathsf{Unif}(0,1),$$

and the A-step yields the sets

$$\Psi_y(u) = \{\psi : F_{t,\psi}(y_1 - 1) \leq u < F_{t,\psi}(y_1)\}, \quad u \in (0,1).$$

For predicting the value of this uniform auxiliary variable, a reasonable choice of predictive random set is the "default" (Martin and Liu 2013)

$$\mathcal{S} = \big[0.5 - |U - 0.5|,\, 0.5 + |U - 0.5|\big], \quad U \sim \mathsf{Unif}(0,1). \tag{13}$$

Then the C-step combines $\Psi_y(\cdot)$ and $\mathcal{S}$ to get $\Psi_y(\mathcal{S}) = \bigcup_{u \in \mathcal{S}} \Psi_y(u)$. Since

$$\Psi_y(\mathcal{S}) \not\ni \psi \iff F_{t,\psi}(y_1 - 1) > \sup \mathcal{S} \text{ or } F_{t,\psi}(y_1) < \inf \mathcal{S},$$

we find that the corresponding plausibility function for singleton ψ is

$$\begin{aligned}
\mathsf{pl}_y(\psi) &= \mathsf{P}_{\mathcal{S}}\{\Psi_y(\mathcal{S}) \ni \psi\} \\
&= 1 - \mathsf{P}_U\{0.5 + |U - 0.5| < F_{t,\psi}(y_1 - 1) \text{ or } 0.5 - |U - 0.5| > F_{t,\psi}(y_1)\} \\
&= 1 - \mathsf{P}_U\{|2U - 1| < 2F_{t,\psi}(y_1 - 1) - 1\} - \mathsf{P}_U\{|2U - 1| < 1 - F_{t,\psi}(y_1)\} \\
&= 1 - \{2F_{t,\psi}(y_1 - 1) - 1\}^+ - \{1 - 2F_{t,\psi}(y_1)\}^+,
\end{aligned}$$

where the "+" superscript denotes the positive part. The somewhat unusual form of this plausibility function is a result of the discreteness of the conditional distribution. Some similar conditioning arguments are used in Jin et al. (2015) to construct an IM for a different version of this discrete problem.

For illustration, I consider two mortality data sets presented in Table 1 of Normand (1999), namely, Trials 1 and 6. Plausibility function for $\log \psi$ for the two data sets are displayed in Figure 3. Both data sets have a relatively small numbers of events, and the two estimated odds ratios are similar: 2.27 in Trial 1 and 2.73 in Trial 6. However, Trial 6 is an overall larger study, so the plausibility function is much more concentrated than that for Trial 1. The flat peak is a result of the discreteness of the problem. These plausibility function plots look quite different than those in Figure 2 Hannig and Xie (2012), based on p-values from Fisher's exact test, in part because they make a certain correction to try to cancel out the effect of the discreteness.

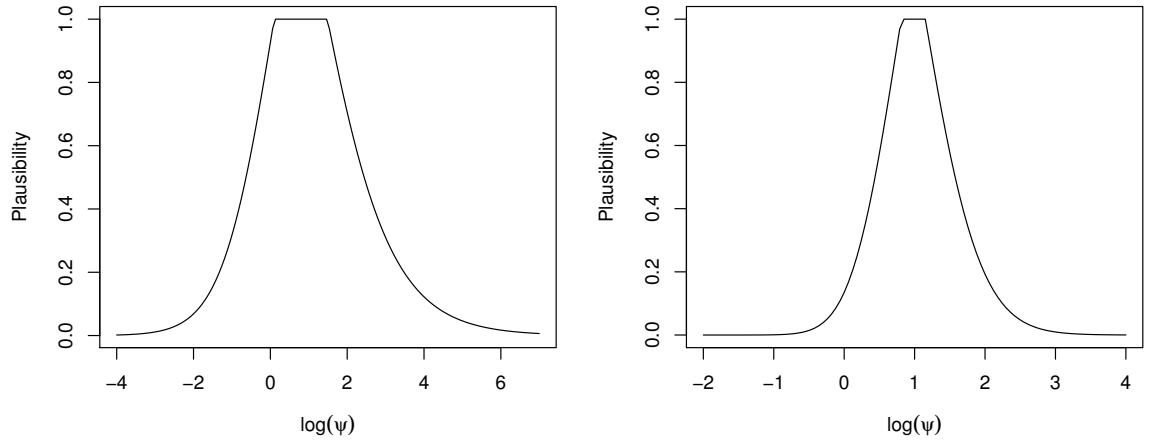

(a) Trial 1: $Y = (1, 2)$, $n = (43, 39)$

(b) Trial 6: $Y = (4, 11)$, $n = (146, 154)$

Figure 3: Plausibility function for the log odds ratio in two mortality data sets (Trial 1 and Trial 6) presented in Table 1 of Normand (1999).

5.2 Error variance in a mixed-effects model

Consider a (possibly unbalanced) normal linear mixed effect model with two variance components, as in Burch and Iyer (1997). The model is written as

$$Y \sim \mathsf{N}_n(X\beta, \sigma_\varepsilon^2 I_n + \sigma_\alpha^2 ZAZ^\top),$$

where X and Z are $n \times p$ and $n \times a$ matrices of predictor variables, respectively, and the parameter is $\theta = (\beta, \sigma_\alpha^2, \sigma_\varepsilon^2)$, with β being the fixed-effect regression coefficients and σ_α^2 and σ_ε^2 the random-effects and error variances, respectively. Suppose that $\psi = \sigma_\varepsilon^2$ is the parameter of interest and $\lambda = (\beta, \sigma_\alpha^2/\sigma_\varepsilon^2)$ is a nuisance parameter. Inference on ψ is interesting from a theoretical point of view because, to my knowledge, there is no method available that can do this exactly; see, also, E et al. (2008, p. 855). In what follows, I also assume that X has full rank $p < n$ and that the matrix A, which describes the correlation structure in the random effects, is known.

Following the setup in E et al. (2008), let K be a $n \times (n-p)$ matrix such that $KK^\top = I_n - X(X^\top X)^{-1}X^\top$ and $K^\top K = I_{n-p}$. It follows that

$$K^\top Y \sim \mathsf{N}_{n-p}(0, \sigma_\varepsilon^2 I_{n-p} + \sigma_\alpha^2 G),$$

where $G = K^\top ZAZ^\top K$ is $(n-p) \times (n-p)$. Let $e_1 > \cdots > e_L \geq 0$ denote the (distinct) eigenvalues of G with multiplicities $r_1, \ldots, r_L$, respectively. Let $P = [P_1, \ldots, P_L]$ be a $(n-p)\times(n-p)$ orthogonal matrix such that $P^\top GP$ is diagonal with eigenvalues $e_1, \ldots, e_L$, in their multiplicities, on the diagonal. For P_ℓ, a $(n-p) \times r_\ell$ matrix, define

$$S_\ell = Y^\top K P_\ell P_\ell^\top K^\top Y, \quad \ell = 1, \ldots, L.$$

Olsen et al. (1976) showed that $(S_1, \ldots, S_L)$ is a minimal sufficient statistic for $(\sigma_\alpha^2, \sigma_\varepsilon^2)$ and, moreover, its distribution is characterized by the equations

$$S_\ell = (\sigma_\alpha^2 e_\ell + \sigma_\varepsilon^2)V_\ell, \quad \ell = 1, \ldots, L,$$

where $V_1, \ldots, V_L$ are independent with $V_\ell \sim \mathsf{ChiSq}(r_\ell)$. By making the transformation from Y to $(S_1, \ldots, S_L)$, the nuisance fixed-effect parameter has been eliminated, as discussed in Section 4.2. With a slight abuse of notation, let $\lambda = \sigma_\alpha^2/\sigma_\varepsilon^2$ be the remaining nuisance parameter. Then the above equation can be rewritten as

$$S_\ell = \psi(\lambda e_\ell + 1)V_\ell, \quad \ell = 1, \ldots, L.$$

In what follows, I propose a generalized IM for ψ using some specialized tricks to eliminate the dependence on λ as much as possible before full marginalization via optimization as discussed in Section 4.2.

Let $\mathcal{L}$ be a proper subset of $\{1, 2, \ldots, L\}$, and write $H(\cdot \mid \lambda) = H_{\mathcal{L}}(\cdot \mid \lambda)$ for the distribution function of $\sum_{\ell \in \mathcal{L}} V_\ell(\lambda)$, a linear combination of independent chi-squares; here, $V_\ell(\lambda) \equiv (\lambda e_\ell + 1)V_\ell$. Next, let $\hat{\lambda}(\cdot)$ be the function that defines maximum likelihood estimator of λ based on observations from the distribution of $V_{-\mathcal{L}}(\lambda)$; like in the R software, the negative subscript means those indices are removed. Define

$$T_{Y,\psi} = H\Big(\frac{1}{\psi}\sum_{\ell \in \mathcal{L}} S_\ell \;\Big|\; \hat{\lambda}(S_{-\mathcal{L}}/\psi)\Big) \tag{14}$$

and

$$Z = H\Big(\sum_{\ell \in \mathcal{L}} V_\ell(\lambda) \;\Big|\; \hat{\lambda}(V_{-\mathcal{L}}(\lambda))\Big) \tag{15}$$

and consider the generalized association

$$T_{Y,\psi} = F_\lambda^{-1}(U), \quad U \sim \mathsf{Unif}(0,1), \tag{16}$$

where F_λ is the distribution function of Z in (15). Note that if $\hat{\lambda}(V_{-\mathcal{L}}(\lambda))$ were exactly equal to λ, then Z would be $\mathsf{Unif}(0,1)$, and the problematic dependence on the nuisance parameter λ would be eliminated. However, it is too much to expect that $\hat{\lambda}(\cdot)$ will exactly equal λ, so the dependence on λ remains, at least for small samples. For the generalized association (16), an appropriate predictive random set for U is the "default" $\mathcal{S}$ used above. Then the construction of the generalized IM for (ψ, λ) is straightforward. Elimination of λ will be carried out by optimizing over λ as discussed in Section 4.2.

For illustration, I will revisit an example presented in Burch and Iyer (1997, Section 4.1) and E et al. (2008, Section 5.2), where $L = 165$, the e's range from $e_1 = 8.56$ to $e_L = 0.57$, and each $r_\ell = 1$ except $r_{105} = 2$. Following Burch and Iyer (1997), I take $\mathcal{L} = \{82, \ldots, 165\}$. Figure 4(a) shows plots of the distribution function F_λ for $\lambda \in \{0.1, 1, 10, 100\}$. This shows that the distribution depends on λ, but maybe not too much. A marginal plausibility interval for ψ, based on this generalized IM, can be obtained by setting $G(\psi) \equiv T_{Y,\psi}$ equal to each of the extreme 2.5% quantiles—optimized over λ—and solving for the corresponding ψ; see, also, Xiong (2015). A plot of $G(\psi)$ for these data is shown in Figure 4(b). In this case, the 95% plausibility interval for ψ is $(0, 3.22)$, which is similar to, but shorter than, the fiducial interval given in E et al. (2008). To check the claimed validity, 2000 independent data sets are simulated by plugging in the maximum likelihood estimator of (ψ, λ). The coverage probability of the 95% marginal plausibility interval is 0.947 and the average length is 3.31.

The fiducial interval being compared to is of high quality (E et al. 2008), so the fact that this generalized IM approach is competitive is quite promising. Theoretically, the

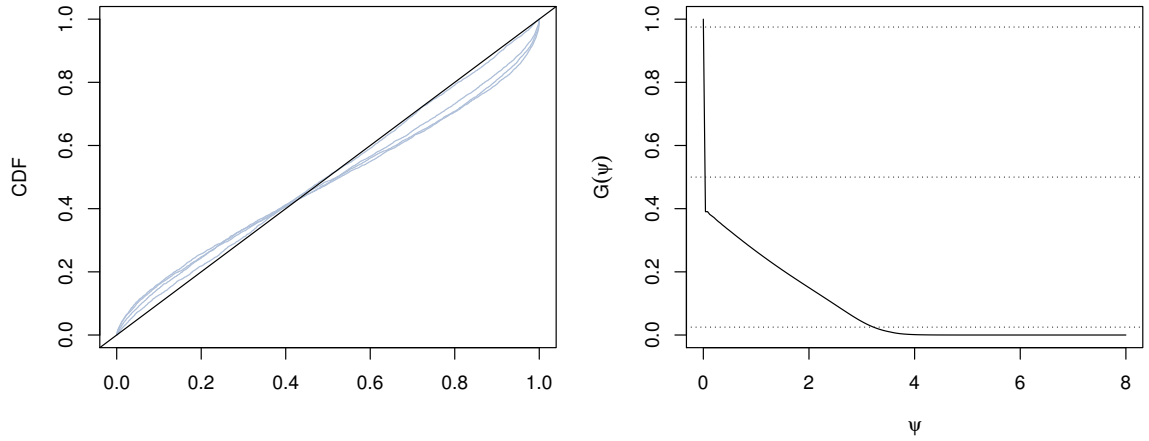

(a) Distribution function F_λ of Z in (15)

(b) Function $G(\psi) = T_{Y,\psi}$ in (14)

Figure 4: Functions related to (14) and (15), with $\mathcal{L} = \{82, \ldots, 165\}$ and the data in Burch and Iyer (1997, Section 4.1); in Panel (a), λ ranges over $\{0.1, 1, 10, 100\}$.

validity result holds, but computation is still a challenge. For one thing, the method of Imhof (1961) used to evaluate $G(\psi)$, as implemented in the `CompQuadForm` package in R, is a bit unstable when ψ is close to zero.

6 Discussion

Previous work on IMs might give the impression that the approach is rigid in its dependence on a version of the data-generating process and, overall, not user-friendly. In this paper, I have proposed a generalized version of the IM framework that is more flexible in a variety of ways. In particular, it makes the IM approach more accessible by seamlessly incorporating some of the more familiar ideas from classical statistics. This added flexibility does not require a sacrifice in terms of the IM's general validity property and, moreover, at least in certain cases, it leads to improved efficiency.

There are at least two important questions that remain to be addressed. First, what is an "optimal" choice of the mapping $T_{y,\theta}$? Some simple ideas were presented in Remark 3 but more work is needed. Second, does this proposed strategy that collapses the problem down to one involving a scalar auxiliary variable work well even in high-dimensional problems? It is likely that this extreme of dimension-reduction will result in a loss of efficiency when the problem is sufficiently complex, but this has yet to be investigated.

Acknowledgments

The author thanks the Guest Editors, Professors Nils Hjort and Tore Schweder, for the opportunity to submit a paper for this special *JSPI* issue, and to Professor Chuanhai Liu and the anonymous reviewers for some very helpful comments on a previous draft.

References

Aickin, M. (2000). Connecting Dempster–Shafer belief functions with likelihood-based inferences. *Synthese*, 123(3):347–364.

Barnard, G. A. (1995). Pivotal models and the fiducial argument. *Int. Statist. Rev.*, 63(3):309–323.

Barndorff-Nielsen, O. E. (1988). *Parametric Statistical Models and Likelihood*, volume 50 of *Lecture Notes in Statistics*. Springer-Verlag, New York.

Barndorff-Nielsen, O. E. (1991). Modified signed log likelihood ratio. *Biometrika*, 78(3):557–563.

Berger, J. O., Bernardo, J. M., and Sun, D. (2009). The formal definition of reference priors. *Ann. Statist.*, 37(2):905–938.

Berger, J. O., Bernardo, J. M., and Sun, D. (2015). Overall objective priors. *Bayesian Anal.*, 10(1):189–221.

Bernardo, J.-M. (1979). Reference posterior distributions for Bayesian inference. *J. Roy. Statist. Soc. Ser. B*, 41:113–147.

Bølviken, E. and Skovlund, E. (1996). Confidence intervals from Monte Carlo tests. *J. Amer. Statist. Assoc.*, 91(435):1071–1078.

Botev, Z. I. and Lloyd, C. J. (2015). Importance accelerated Robbins–Monro recursion with applications to parametric confidence limits. *Electron. J. Stat.*, 9:2058–2075.

Brazzale, A. R., Davison, A. C., and Reid, N. (2007). *Applied Asymptotics: Case Studies in Small-Sample Statistics*. Cambridge University Press, Cambridge.

Burch, B. D. and Iyer, H. K. (1997). Exact confidence intervals for a variance ratio (or heritability) in a mixed linear model. *Biometrics*, 53(4):1318–1333.

Claggett, B., Xie, M., and Tian, L. (2014). Meta-analysis with fixed, unknown, study-specific parameters. *J. Amer. Statist. Assoc.*, 109(508):1660–1671.

Davison, A. C. and Hinkley, D. V. (1997). *Bootstrap Methods and Their Application*, volume 1. Cambridge University Press, Cambridge.

Dawid, A. P. and Stone, M. (1982). The functional-model basis of fiducial inference. *Ann. Statist.*, 10(4):1054–1074. With discussion.

Dempster, A. P. (2008). The Dempster–Shafer calculus for statisticians. *Internat. J. Approx. Reason.*, 48(2):365–377.

Denœux, T. (2014). Likelihood-based belief function: justification and some extensions to low-quality data. *Internat. J. Approx. Reason.*, 55(7):1535–1547.

E, L., Hannig, J., and Iyer, H. (2008). Fiducial intervals for variance components in an unbalanced two-component normal mixed linear model. *J. Amer. Statist. Assoc.*, 103(482):854–865.

Fisher, R. A. (1973). *Statistical Methods and Scientific Inference*. Hafner Press, New York, 3rd edition.

Fraser, D. A. S. (1990). Tail probabilities from observed likelihoods. *Biometrika*, 77(1):65–76.

Fraser, D. A. S. (1991). Statistical inference: likelihood to significance. *J. Amer. Statist. Assoc.*, 86(414):258–265.

Fraser, D. A. S. (2011). Is Bayes posterior just quick and dirty confidence? *Statist. Sci.*, 26(3):299–316.

Garthwaite, P. H. and Buckland, S. T. (1992). Generating Monte Carlo confidence intervals by the Robbins–Monro process. *J. Roy. Statist. Soc. Ser. C*, 41(1):159–171.

Ghosh, M. (2011). Objective priors: An introduction for frequentists. *Statist. Sci.*, 26(2):187–202.

Hannig, J. (2009). On generalized fiducial inference. *Statist. Sinica*, 19(2):491–544.

Hannig, J., Iyer, H., Lai, R. C. S., and Lee, T. C. M. (2015). Generalized fiducial inference: A review. *J. Amer. Statist. Assoc.*, to appear.

Hannig, J. and Xie, M.-g. (2012). A note on Dempster–Shafer recombination of confidence distributions. *Electron. J. Stat.*, 6:1943–1966.

Imhof, J. P. (1961). Computing the distribution of quadratic forms in normal variables. *Biometrika*, 48:419–426.

Jin, H., Li, S., and Jin, Y. (2015). The IM-based method for testing non-inferiority of odds ratio in matched-pairs design. *Unpublished manuscript*.

Liu, D., Liu, R. Y., and Xie, M. (2015). Multivariate meta-analysis of heterogeneous studies using only summary statistics: efficiency and robustness. *J. Amer. Statist. Assoc.*, 110(509):326–340.

Liu, D., Liu, R. Y., and Xie, M.-g. (2014). Exact meta-analysis approach for discrete data and its application to 2×2 tables with rare events. *J. Amer. Statist. Assoc.*, 109(508):1450–1465.

Martin, R. (2015). Plausibility functions and exact frequentist inference. *J. Amer. Statist. Assoc.*, 110:1552–1561.

Martin, R. and Liu, C. (2013). Inferential models: A framework for prior-free posterior probabilistic inference. *J. Amer. Statist. Assoc.*, 108(501):301–313.

Martin, R. and Liu, C. (2015a). Conditional inferential models: Combining information for prior-free probabilistic inference. *J. R. Stat. Soc. Ser. B*, 77(1):195–217.

Martin, R. and Liu, C. (2015b). *Inferential Models: Reasoning with Uncertainty*. Monographs in Statistics and Applied Probability Series. Chapman & Hall/CRC Press.

Martin, R. and Liu, C. (2015c). Marginal inferential models: Prior-free probabilistic inference on interest parameters. *J. Amer. Statist. Assoc.*, 110:1621–1631.

Normand, S. (1999). Meta-analysis: formulating, evaluating, combining, and reporting. *Stat. Med.*, 18(3):321–359.

Olsen, A., Seely, J., and Birkes, D. (1976). Invariant quadratic unbiased estimation for two variance components. *Ann. Statist.*, 4(5):878–890.

Pal Majumdar, A. and Hannig, J. (2015). Higher order asymptotics of generalized fiducial distributions. *Unpublished manuscript*.

Reid, N. (2003). Asymptotics and the theory of inference. *Ann. Statist.*, 31(6):1695–1731.

Robbins, H. and Monro, S. (1951). A stochastic approximation method. *Ann. Math. Statistics*, 22:400–407.

Schweder, T. and Hjort, N. L. (2002). Confidence and likelihood. *Scand. J. Statist.*, 29(2):309–332.

Schweder, T. and Hjort, N. L. (2016). *Confidence, Likelihood, Probability: Statistical Inference with Confidence Distributions*. Cambridge Univ. Press.

Shafer, G. (1987). Belief functions and possibility measures. In Bezdek, J. C., editor, *The Analysis of Fuzzy Information, Vol. 1: Mathematics and Logic*, pages 51–84. CRC.

Singh, K., Xie, M., and Strawderman, W. E. (2007). Confidence distribution (CD)—distribution estimator of a parameter. In *Complex datasets and inverse problems*, volume 54 of *IMS Lecture Notes Monogr. Ser.*, pages 132–150. Inst. Math. Statist., Beachwood, OH.

Taraldsen, G. and Lindqvist, B. H. (2013). Fiducial theory and optimal inference. *Ann. Statist.*, 41(1):323–341.

Wasserman, L. A. (1990). Belief functions and statistical inference. *Canad. J. Statist.*, 18(3):183–196.

Xie, M. and Singh, K. (2013). Confidence distribution, the frequentist distribution of a parameter – a review. *Int. Statist. Rev.*, 81(1):3–39.

Xie, M., Singh, K., and Strawderman, W. E. (2011). Confidence distributions and a unifying framework for meta-analysis. *J. Amer. Statist. Assoc.*, 106(493):320–333.

Xiong, S. (2015). Local optimization-based statistical inference. Unpublished manuscript, `arXiv:1502.00465`.

Yang, G., Liu, D., Liu, R. Y., Xie, M., and Hoaglin, D. C. (2014). Efficient network meta-analysis: a confidence distribution approach. *Stat. Methodol.*, 20:105–125.